

## Halftone Watermarking and Related Applications

### Related Application Data

This patent application is a continuation of U.S. Application No. 09/689,226, filed October 11, 2000 (now U.S. Patent No. 6,694,041), which is hereby incorporated by reference.

### Technical Field

The invention relates to multimedia signal processing, and in particular relates to image watermarking methods and related applications.

### Background and Summary

Digital watermarking is a process for modifying physical or electronic media to embed a machine-readable code into the media. The media may be modified such that the embedded code is imperceptible or nearly imperceptible to the user, yet may be detected through an automated detection process. Most commonly, digital watermarking is applied to media signals such as images, audio signals, and video signals. However, it may also be applied to other types of media objects, including documents (e.g., through line, word or character shifting), software, multi-dimensional graphics models, and surface textures of objects.

Digital watermarking systems typically have two primary components: an encoder that embeds the watermark in a host media signal, and a decoder that detects and reads the embedded watermark from a signal suspected of containing a watermark (a suspect signal). The encoder embeds a watermark by altering the host media signal. The reading component analyzes a suspect signal to detect whether a watermark is present. In applications where the watermark encodes information, the reader extracts this information from the detected watermark.

Several particular watermarking techniques have been developed. The reader is presumed to be familiar with the literature in this field. Particular techniques for

embedding and detecting imperceptible watermarks in media signals are detailed in the assignee's US Patent Nos. 6,614,914 and 5,862,260, which are hereby incorporated by reference.

This disclosure describes halftone image watermark methods and systems. One aspect of the disclosure is a method of halftone image watermarking. This method assigns a set of halftone watermark values to locations within a halftone image. It diffuses error associated with the halftone watermark values to neighboring locations within the halftone image. The error is characterized as a difference between a multilevel pixel value at a location of a halftone watermark dot, and the halftone watermark value of the halftone watermark dot. This method may be used in conjunction with other watermark embedding stages. For example, a robust watermark carrying a key for the halftone watermark may be embedded in the image before embedding the halftone watermark.

Another aspect of the disclosure is a method of decoding a halftone watermark from an image. The decoding method uses a key to identify locations and values of a halftone watermark, and analyzes pixel values at the locations to determine whether the values at the locations correspond to the values specified by the key. In one application, a robust watermark embedded in the image carries the key used to decode the halftone watermark from the image. A decoder reads the robust watermark from a scan of the halftone watermarked image, optionally compensating for geometric distortion of the scanned image using an orientation signal. The decoder then extracts the key from the robust watermark and passes it to a verifier, which in turn, examines a high resolution scan of the suspect image to determine whether the halftone watermark is present at locations specified by the key.

Another aspect of the disclosure is another method of embedding a watermark in a halftone image. This method computes a watermark image comprising an array of values corresponding to pixel locations in a halftone image. It embeds the watermark image in the halftone image by using the values of the watermark image to modulate thresholds at the pixel locations. The thresholds are used in a halftone process to convert

multilevel pixel values in a multilevel per pixel image into halftone pixel values of the halftone image.

Another aspect of the disclosure is yet another method of embedding a watermark in a halftone image. This method computes a watermark image comprising an array of values corresponding to pixel locations in a target halftone image at a halftone resolution. It then combines the array of values of the watermark image with corresponding multilevel pixel values of a multilevel per pixel image to create a watermarked multilevel per pixel image at the resolution of the target halftone image. Finally, it performs a halftone process to convert the watermarked multilevel per pixel image to a watermarked halftone image.

Finally, another aspect of the disclosure is a watermark decoder. The decoder includes a watermark detector and reader. The detector analyzes portions of an image to detect a watermark signal embedded in the image. The image is scanned from a halftone printed image at a sufficiently high resolution to discern a watermark image embedded at a resolution of the halftone image. The watermark reader reads a watermark signal from the portions of the image and decodes an auxiliary message comprising one or more symbols.

Further features will become apparent with reference to the following detailed description and accompanying drawings.

20

## Brief Description of the Drawings

Fig. 1 is a diagram illustrating an example of an error diffusion mask used in creating halftone images from multilevel per pixel images.

Fig. 2 is a diagram illustrating another example of an error diffusion mask used in creating halftone images from multilevel per pixel images.

Fig. 3 is a diagram illustrating another example of an error diffusion mask used in creating halftone images from multilevel per pixel images.

Fig. 4 is a diagram of an application of a halftone watermark application.

Fig. 5 is a diagram illustrating a method of creating a watermarked halftone image using a threshold mask.

Fig. 6 is a diagram illustrating a method using the halftone screen as an orientation signal to determine geometric distortion of a watermarked image signal.

5

## Detailed Description

The following description details methods for watermarking halftone images and related applications. While the watermarking methods apply to other forms of halftoning, the description provides specific examples applicable to error diffusion techniques used to create halftone images.

10 One form of error diffusion used in processing halftone images is called Floyd-Steinberg error diffusion. In a typical implementation for 8 bit per pixel images, this error diffusion method takes a 0-255 level input image at a first resolution (e.g., 200 pixels per inch) and produces a binary image with a higher resolution (e.g., 600 dots per inch). This description details an implementation for an image plane where each pixel  
15 has a corresponding multilevel value, such as a luminance value, or some other color channel like cyan, magenta, yellow, etc. The description applies to color images with more than one multilevel value per pixel. In such cases, a halftone process operates on each of the color channels per pixel and creates a halftone image for each channel.

As a first step, the input image is upsampled to the higher resolution. One  
20 approach for producing the binary 600 dots per inch (dpi) image from the upsampled 0-255 level 600 pixel per inch (ppi) image is to threshold the pixel values so that a pixel value less than 128 produces a zero in the corresponding location of the binary image; otherwise the corresponding location would be set to a one. Error diffusion takes this general approach in a raster-scan order, but achieves improved performance by  
25 "diffusing" the error at each location to nearby locations yet to be processed. This error diffusion is guided by a weighting mask, shown in Fig. 1.

The "X" represents the pixel location currently being processed, and the adjoining cells show how the error from that location is diffused. For this algorithm, the sum of the diffused errors is equal to the total error at location X. For any specific binary location,

the accuracy of this method is no better than the simple thresholding approach. However, if a pixel from the original 200 ppi image is compared with the corresponding 3x3 binary region, the number of cells with a one is closely correlated with the multilevel pixel value.

5 A more precise description of the basic error diffusion algorithm uses three equations. See "Digital Color Halftoning", p. 359, by Kang (co-published by The International Society for Optical Engineering and IEEE Press 1999), which is incorporated by reference:

$(m, n)$ : binary (upsampled) pixel location

$p_i(m, n)$ : upsampled input intensity, 0 - 255

$p'(m, n)$ : modified input intensity

$p_o(m, n)$ : output binary value

$e(m, n)$ : error at location  $(m, n)$

$w_{kl}$ : error diffusion weighting coefficients

$$10 \quad p'(m, n) = p_i(m, n) + \sum_{kl} w_{kl} e(m - k, n - l)$$

$$p_o(m, n) = \begin{cases} 1 & : p'(m, n) \geq 128 \\ 0 & : p'(m, n) < 128 \end{cases}$$

$$e(m, n) = p'(m, n) - 255 p_o(m, n)$$

The weighting mask shown in Fig. 1 can be modified to show which errors are diffused into a given location. Fig. 2 shows an example of such modifications.

15 In the following sections, we describe a modified error diffusion method that embeds a watermark comprising a set of binary values at specified dot locations in a binary image. This method starts with an upsampled binary host image, a list of dot locations in a binary image and corresponding binary values for a watermark. This method assigns to these locations the corresponding values of the watermark and tries to

improve the image with an error diffusion algorithm at the other (non-watermark) locations. It is assumed that the fraction of total pixels occupied by the watermark is small.

In the error diffusion algorithm introduced above, the error for a given location is  
5 always calculated after the algorithm has processed all preceding locations. Typically,  
the algorithm scans a rectangular image comprising scanline rows of pixels starting from  
the top row and scanning from left to right across each row. The modified error diffusion  
method calculates error values for all locations at the start and modifies them as the  
method proceeds. At the start, the error of all locations not covered by the watermark are  
10 set to zero. For a location covered by the watermark, the error is calculated as

$$e(m, n) = p_i(m, n) - 255W(m, n)$$

where  $W(m, n)$  is the value of the watermark at location  $(m, n)$ . A new set of error  
15 diffusion weights is used; the pattern of error diffusion into location X is shown in Fig. 3.

In this method, the diffusion of errors takes place in two directions. Errors from  
the watermarked locations are diffused in one direction (backwards in this case), and  
errors from processed locations are diffused in another direction (forwards). As each  
location is processed, the error for that location is updated for later diffusion. The  
20 calculation of the output binary value is unchanged from the error diffusion algorithm,  
with the exception at watermark locations, where  $p_o(m, n)$  is set to  $W(m, n)$ .

The appearance of the watermark can be improved by arranging locations of  
watermark dots in a pseudorandom pattern subject to some additional human visual  
system frequency response criteria. One approach for arranging the locations of the  
25 watermark dots is to use a minimum visual cost technique. In this approach the  
watermark is chosen by minimizing a visual cost function such as

$$C = \iint |H(f_x, f_y)V(f_x, f_y)|^2 df_x df_y$$

where  $H(f_x, f_y)$  is the frequency spectrum of the watermark and  $V(f_x, f_y)$  is a frequency response model of the human visual system. One possible such model is due to Sullivan et al. and given in Kang, section 5.6. The effect of this cost function is to weight the frequency content of the watermark by the human visual ability to detect it. In 5 this manner, watermark patterns that move the frequency content to a region of the frequency spectrum where the human visual system is less sensitive will have a lower cost. This method of selecting locations and values of watermark dots provides a perceptually more even distribution of the dots in solid regions (e.g., white regions in a grayscale image). Searching for a watermark with a low cost can be done by a variety of 10 methods; one such method is simulated annealing.

The encoder may repeat the watermark in blocks of the host image, such as contiguous blocks of pixels throughout the image. Each block may use the same or a different key. If the encoder uses different keys per block, each key may be related to another key by a secret function, such as a cryptographic function.

15 To read the watermark, a watermark decoder uses a key specifying where the watermark dots are located. It then determines the binary values of those dots. In cases where the watermark is embedded in a printed image, a scanner or camera with sufficiently high resolution to read the halftone dots creates a digital image from which the decoder extracts the watermark.

20 There are many applications of this type of watermark. It may be used to carry a message, including usage control instructions or other metadata like an identifier of the image owner or an index to a database record storing related information. It may also convey a fixed pattern used to determine whether a watermarked image (e.g., a printed halftone watermarked image) is authentic or has been altered. In such an application, the 25 decoder compares the extracted watermark pattern with the known pattern, and based on this comparison determines whether the watermarked image has been altered (e.g., copied, scanned and re-printed, compressed, etc.). It can also specify where the image has been altered by showing locations in the scanned image where the halftone watermark is not present.

Since the watermark is applied to halftone images it may be applied as part of the printing process where a multilevel per pixel digital image is converted to a halftone image. For example, it may be incorporated into the halftoning process implemented in a printer device or printer driver executing in a computer that sends the image to the printer for printing. More generally, the watermarking method may be applied in applications where halftone images are printed, including commercial printing presses as well as personal printing devices, such as ink jet printers.

In some types of image watermarking, the host image interferes with the watermark. The host image provides a communication channel for the watermark. In decoding the watermark from a host image, the host image can be considered noise that interferes with the ability of the watermark decoder to accurately extract the watermark signal. To increase the chances of accurate recovery of the watermark, it is spread throughout a large portion (perhaps the entire image) of the host image.

The type of watermark described in the previous paragraph can be embedded directly into a halftone image. The following discussion details an error diffusion method for directly embedding this type of watermark into a halftone image.

In this method, the watermark to be embedded,  $W(m, n)$ , takes values from 0-255. The error diffusion process is changed so that the threshold used to calculate the binary output values is modulated by the watermark signal:

20

$$p_o(m, n) = \begin{cases} 1 & : p'(m, n) \geq T(m, n) \\ 0 & : p'(m, n) < T(m, n) \end{cases}$$

$$T(m, n) = 128 - I(W(m, n) - 128)$$

The intensity level  $I$  controls how heavily the watermark is embedded in the image.

As an alternative to modulating error diffusion thresholds, the watermark may be embedded without modifying the halftoning process. For example, a multilevel per pixel watermark signal is created at the resolution of a target halftone image. The watermark encoder produces the multilevel per pixel watermark signal at the desired resolution of

the halftone image, or at some other resolution and up or down samples it to match the resolution of a target halftone image. This watermarked signal is then added to the host image at the same spatial resolution to create a composite, watermarked image. The error diffusion process or some other type of halftone process may then be applied directly to  
5 this composite image to generate a watermarked halftone image. This technique applies to a variety of halftone processes including ordered dithering (e.g., blue noise masks, clustered dot halftones, etc.) as well as error diffusion halftone processes.

There are a variety of ways to generate the watermark signal. One approach is to take an auxiliary message comprising binary or M-ary symbols, apply error correction  
10 coding to it, and then spread spectrum modulate the error correction encoded message.

One way to spread spectrum modulate the message is to spread each binary symbol in the message over a pseudorandom number, using an exclusive OR operation or multiplication operation. The resulting binary message elements in the spread spectrum modulated message signal are then mapped to spatial image locations. The watermark  
15 signal may be expressed in a binary antipodal form, where binary symbols are either positive or negative. To increase robustness, the spread spectrum modulated message signal may be repeated throughout the host image, by for example, embedding the message signal in several blocks of the host image. In particular, the watermark encoder may embed instances of the watermark signal into contiguous blocks of pixels throughout  
20 a portion of the host image or throughout the entire host image.

Perceptual modeling may be applied to the host image to calculate a gain vector with gain values that correspond to the message signal elements. For example, in the case where the upsampled watermarked signal is added to the host signal, the gain values may be used to scale binary antipodal values of the message signal before adding them to  
25 the host signal. Each gain value may be a function of desired watermark visibility and detectability constraints. In particular, the perceptual model analyzes the image to determine the extent to which it can hide a corresponding element of the watermark image. One type of an analysis is to compute local contrast in a neighborhood around each pixel (e.g., signal activity) and select a gain for a pixel as a function of local

contrast. A detectability model analyzes the host signal to determine the extent to which pixel values are biased toward the value of the watermark signal at the corresponding pixel locations. It then adjusts the gain up or down depending on the extent to which the host image pixels are biased towards the watermark signal.

5 This type of watermark may be read from the watermarked halftone image or other image representations of that watermarked image, such as a multilevel per pixel representation of the image at a resolution sufficiently high to represent the watermark signal. To decode the watermark, a watermark decoder detects the presence and orientation of the watermark in the watermarked image. It then performs an inverse of  
10 the embedding function to extract an estimate watermark message signal.

The message signal is robustly encoded using a combination of the following processes:

1. repetitively encoding instances of a message signal at several locations (e.g., blocks of the image);
- 15 2. spread spectrum modulation of the message, including modulation techniques using M sequences and gold codes; and
3. error correction coding, such as convolution coding, turbo coding, BCH coding, Reed Solomon coding, etc.

The watermark decoder reconstructs an embedded message from the estimated  
20 watermark signal by:

1. aggregating estimates of the same message element in repetitively encoded instances of the message;
2. performing spread spectrum demodulation, and
3. error correction decoding.

25 In one implementation, the decoder uses an orientation signal component of the watermark to detect its presence and orientation in the watermarked image. It then performs a predictive filtering on the image sample values to estimate the original un-watermarked signal, and subtracts the estimate of the original from the watermarked signal to produce an estimate of the watermark signal. It performs spread spectrum

demodulation and error correction decoding to reconstruct an auxiliary message embedded in the watermarked signal.

For more details about embedding an image watermark, and detecting and reading the watermark from a digitized version of the image after printing and scanning see

5 assignee's US Patent Nos. 6,614,914 and 5,862,260, which are hereby incorporated by reference. In order to make the watermark robust to geometric distortion, the watermark includes an orientation watermark signal component. Together, the watermark message signal and the orientation watermark signal form the watermark signal. Both of these components may be added to a host image at the resolution of the halftone image before  
10 the host image is converted to a the halftone image. Alternatively, these components may be combined to form the watermark signal used in modulating the error diffusion threshold used in an error diffusion type halftone process.

One type of watermark orientation signal is an image signal that comprises a set of impulse functions in the Fourier magnitude domain, each with pseudorandom phase.

15 To detect rotation and scale of the watermarked image (e.g., after printing and scanning of the watermarked image), the watermark decoder converts the image to the Fourier magnitude domain and then performs a log polar resampling of the Fourier magnitude image. A generalized matched filter correlates the known orientation signal with the re-sampled watermarked signal to find the rotation and scale parameters providing the highest correlation. The watermark decoder performs additional correlation operations  
20 between the phase information of the known orientation signal and the watermarked signal to determine translation parameters, which identify the origin of the watermark message signal. Having determined the rotation, scale and translation of the watermark signal, the reader then adjusts the image data to compensate for this distortion, and  
25 extracts the watermark message signal as described above.

The halftone watermarks described above may be used in combination with one or more other watermarks. In one application, for example, a robust watermark is used to carry a key that specifies the dot locations of a halftone watermark. In particular, the robust watermark's message payload carries a key that identifies specific dots (the high-

resolution binary values) that were turned on or off in a specific pattern. These binary valued bits act as a secondary fragile watermark that can be verified by close inspection of the image.

Fig. 4 illustrates an implementation of this application. On the watermark embedding side, a watermark encoder 100 operates on an input image 102 to embed a robust watermark that survives printing, scanning, and geometric distortion. An example of this type of watermark is described above and in the patent and patent application incorporated by reference. At least part of the message payload of this robust watermark carries a key that is used to decode a halftone watermark at specified halftone dot locations in a halftone image.

In this implementation, the key includes a seed to a random number generator 104 that identifies locations of the halftone watermark dots in the halftone image and also specifies the binary values (on or off) of the dots at those locations. A halftone converter 106 (e.g., error diffusion, blue noise masking, etc.) then converts the robustly watermarked image into a halftone image while ensuring that the halftone watermark dots are assigned the correct values based on the output of the random number generator.

One example of this process is the modified error diffusion method described above in which the halftone converter diffuses the error introduced by halftone watermark dots to a spatial neighborhood around their respective locations while ensuring that those watermark dot values remain fixed. This approach reduces the perceptible impact of the halftone watermark. Other halftone methods may be used as well as long as they ensure that the halftone watermark dots are set as specified by the key. The result is a halftone image 108 that may be printed using conventional printer technology, such as ink jet printing, etc. (110). In fact, the watermark embedding process (100) may be implemented in a printer or software driver for a printer.

To verify the authenticity of the printed image, a scanner 120 captures a digital image from a printed image 122. Since the robust watermark survives printing and then scanning by a low resolution (e.g., 100 dpi) scanner or digital camera, it may be recovered from a digital image captured from the printed image 122. A watermark reader

124, using the decoding operations outlined above, detects the robust watermark, determines its orientation, and then reads the watermark payload, including the key. It supplies the key to a random number generator 126, which provides the halftone dot locations and values of the halftone watermark.

5        A secondary watermark verifier 128 then examines a suitably high resolution scan of the printed image to determine whether the halftone watermark is present. A “suitably high resolution scan” of the printed image is one in which the halftone dots of the printed image are readable. This high resolution scan may be the same image captured by the scanner 120 if it is at a sufficient resolution. Alternatively, a separate image capture  
10      device 130 may be used to capture a high resolution image depicting the halftone dots. The verifier 128 provides a signal indicating the extent to which the halftone watermark is present.

15      Based on the output of the verifier, a number of actions can be taken. Some actions include recording identifying information about the user (e.g., a device ID of the user’s computer or imaging device, address (e.g., network address), user ID, etc.), sending this information and the result of the verification operation to a remote device via communication link, displaying information about rules governing use of the image, connecting the decoding system to a licensing or electronic transaction server (e.g., a web server) enabling the user to license or purchase related rights, products or services, etc.  
20      electronically.

Both the robust and fragile halftone watermark may also carry other information. They may be used to carry a message, including usage control instructions or other metadata like an identifier of the image owner, a computer address for establishing a remote connection (e.g., an IP address, URL, etc.) or an index to a database record  
25      storing related information. In one application, the message carries an identifier that is used to fetch related information to the image. In particular, the watermark decoder communicates the identifier to a database management system in the form of a request. The database management system and underlying database records may be implemented in a remote device connected to the watermark decoder device via a network connection

(e.g., a web server on the Internet connected via a TCP/IP connection) or within the system containing the watermark decoder (e.g., a local database executing within the same device as the decoder or within a computer locally connected to the decoding device via a port, such as a serial, parallel, infrared, Bluetooth wireless or other peripheral port). In response to the request, the database looks up information related to the identifier or identifiers extracted from the robust watermark. The database either returns the information to the watermark decoding device (e.g., a personal computer, personal digital assistant, Internet appliance, scanner, printer, etc.), or forwards it to one or more other devices along with an address of the decoding device, which in turn, return related information to the decoding device. One example is to use an identifier or address encoded in the robust watermark to fetch information from a web site related to the watermarked image. For more information on this use of the robust watermark, see US Patent Nos. 5,841,978, 6,505,160 and 6,411,725, and co-pending application 09/571,422, which are hereby incorporated by reference.

As noted previously, for many image watermark applications, it is necessary to, as part of the watermark detection and reading process, determine the orientation (scale, rotation, etc.) of the watermarked image prior to the process of reading the watermark. One way to determine orientation of a watermarked image relative to its orientation at the time of embedding the watermark is to use some form of calibration signal as part of the watermark. The calibration signal may be integrally related to the watermark message signal, such as a carrier signal or synchronization code. Alternatively, it may be a distinct signal, such as an imperceptible registration template.

The following watermarking method uses a halftone image screen as an orientation signal, avoiding the need to embed a separate calibration signal. The halftone screen may be one normally used to print un-watermarked halftone images, or one specifically adapted to create and print watermarked halftone images.

This method is suitable for image halftoning methods in which the halftoning is implemented through the use of threshold masks. A threshold mask is a pattern which can be tiled in a repetitive fashion over an image. Most commonly, the threshold mask is

a square matrix of numbers. For example, the threshold mask may be 128x128, and the image may be comprised of an array of 8-bit pixels. In this case, the threshold mask contains numbers ranging from 0 to 255, corresponding to the 256 possible levels in an 8 bit pixel value. In cases where the number of elements in the mask exceeds the number  
5 of levels per pixel, the threshold mask is usually designed to include an equal number of occurrences of each level. For example, in the current example where there are 16384 elements in a 128x128 mask, there are 64 entries of the threshold mask for each level ( $16384/256 = 64$ ). It is possible to use a 16x16 mask with each element corresponding to a different level from 0 to 255. However, when tiled across the image, such a mask may  
10 introduce more visual artifacts than the larger mask. Of course, these masks are just examples, and there are many alternative combinations of mask sizes, dimensions, and levels for threshold masks used in halftone screen processes.

In a typical halftone screen process, the threshold mask is used to create the binary pattern of ink dots used to represent the printed version of an image. If an image  
15 with 200 pixels per inch is to be printed using 2400 dots per inch of resolution, the image is first upsampled to 2400 pixels per inch. Then the threshold mask is tiled and applied to the upsampled image to construct the halftone image. Each pixel of the halftone image is calculated from the corresponding pixel of the upsampled image and the value of the tiled threshold mask at that location. If the threshold mask value is less than or equal to the  
20 upsampled image pixel, then the halftone image pixel is set to 1; otherwise it is set to 0.

Error diffusion halftoning may not be implemented in this way.

Fig. 5 is a diagram illustrating a method of creating a watermarked halftone image using a threshold mask. The watermark encoding process begins by transforming an auxiliary message into a watermark image signal (150). The method shown in Fig. 5 can  
25 be used with a variety of methods of constructing a watermark image signal. For the purpose of illustration, Fig. 5 cites an example using a spatial spread spectrum watermark. In this case, a watermark message payload is spread spectrum modulated over an area coextensive with the threshold mask. Alternative watermark encoding functions may be used, such as frequency domain techniques, where the watermark signal

modulates frequency coefficients to encode message symbols, or statistical feature modulation where the watermark modulates features of the signal, such as its autocorrelation, power, amplitude, signal peaks, etc. Each of these embedding functions can be characterized as adding a spatial watermark image signal with corresponding elements in the host image.

The encoder may modulate the watermark image signal with a gain to make the watermark message more likely to be recovered, less visible, or some balancing of recoverability and imperceptibility. In the case of the spatial spread spectrum watermark image, the encoder converts the message to a binary antipodal signal (150) and a gain modulator adjusts the value of each element of that signal (152). The spatial resolution of the watermark image signal corresponds to the target halftone image.

The watermark image signal may be tiled, so that the same payload is repeated over the image, or it may be varied, so that different information is embedded in different areas of the image.

To prepare the host image for the watermark, the encoder takes a multilevel per pixel form of the image (154) and converts it to the halftone resolution (156). This step typically involves upsampling the multilevel image to a higher resolution (e.g., upsampling a 200 dpi, 8 bit per pixel image to a 2400 dpi per pixel image). One method for inserting the watermark is to add a spatial domain representation of the watermark image signal with the multilevel image signal (158), each at the halftone image resolution, to create a composite image.

Next, the encoder applies a threshold mask (160, 162) to the resulting composite image. The halftone screen process converts pixel levels below or equal to the corresponding threshold mask level to zero, otherwise it sets them to one. The result is a watermarked halftone image. The image may then be printed using a variety of halftone printing process, including ink jet printers and printing presses.

Fig. 6 is a diagram illustrating a method using the halftone screen as an orientation signal to determine geometric distortion of a watermarked image signal. In

particular, this method is used to calculate the orientation of the host signal at the time of watermark embedding to facilitate recovery of a watermark message.

First, a scanner or camera or other imaging device captures a digital version of the printed, watermarked image 170 at sufficiently high resolution (172) to distinguish the halftone pixels. In the continuing example, a watermark decoder obtains blocks of a digital image 174, each having an area approximately corresponding to the size of the threshold mask (e.g., 128x128). Each digital image pixel is approximately equal in size to a halftone image pixel.

Next, the decoder filters the image block using a low pass filtering technique 176 like computing the average pixel value and replacing each pixel with the average. Next, the decoder applies the same threshold mask used to create the halftone image for printing (178, 180). This process results in a target orientation signal having orientation attributes similar to the original image.

To determine the geometric distortion of the received image relative to the original, watermarked image, the decoder performs a series of correlation operations between the target orientation signal and the received image. In the method shown in Fig. 6, the decoder converts both the received image block and the target orientation signal derived from it to a correlation domain (182). In particular, it performs a 128x128 Fast Fourier Transform and resamples the resulting Fourier magnitude representation into a log-polar coordinate system to get a Fourier-Mellin representation of both images.

Correlation operators 184, such as a generalized matched filter, correlate the Fourier-Mellin representation of the target signal (the reference signal) with a Fourier-Mellin representation of the received image. The correlation operation produces an estimate of the scale and rotation parameters of the digital image.

The decoder transforms the received spatial digital image by the inverse of the estimated rotation and scale. The correlation operator correlates the resulting spatial image with the spatial target orientation signal to obtain an estimate of the translation of the image with respect to the threshold mask.

Having estimates for the scale, rotation, and translation parameters of the digital version of the watermarked image, a watermark message reader (186) demodulates the spread spectrum watermark signal from the received image.

Each of the halftone watermark methods described above may be used in “fragile” watermark or “semi-fragile” watermark applications, where the watermark is analyzed to detect and characterize alterations to a watermarked image. A fragile watermark refers to a watermark that degrades or becomes unreadable when the host image is subjected to certain types of distortion. A semi-fragile watermark is a variant on this theme where the watermark survives some types of distortion, but not others.

One approach is to design the watermark so that it degrades or becomes unreadable in response to certain types of distortions, like printing, scanning on consumer grade scanners, image compression, etc. If the decoder is unable to detect or read the watermark, or the measured degradation exceeds a threshold, then the decoder provides output indicating that the image has been altered (e.g., is not authentic). The degradation can be measured by the extent of recovery of the watermark signal, such as the extent of correlation between a known watermark signal and one extracted from a received image. In such fragile or semi-fragile applications the gain values of the halftone watermark methods previously described may be chosen to provide an explicit balance between ease of reading and fragility with respect to certain types of distortions.

An enhanced approach is to design the watermark such that the type of alteration can be distinguished. For example, the type of alteration can be characterized by the type of degradation it causes to the watermark signal. By quantifying the degradation to different aspects of the watermark signal, a decoder can match the observed degradation to a particular type of alteration.

The performance of such fragile watermarking applications can be improved by choosing the frequency distribution of the embedded watermark to differentiate between different forms of distortion. There are two main considerations in designing the watermark signal attributes. The first is to put part of the watermark’s energy in frequencies that are likely to be degraded by a particular type of degradation to be

detected (such as print and scan operations); the second is to put energy in the frequencies that are less likely to be degraded by distortion that the watermark should survive (such as smudges, and normal wear and tear). By analyzing the frequency distribution of the watermark, the decoder can distinguish between forms of alteration.

- 5 This approach is particularly useful to determine whether a printed image is genuine (e.g., whether it has been reproduced through a scan and print operation), as opposed to being merely soiled and worn through ordinary use.

### ***Concluding Remarks***

Having described and illustrated the principles of the technology with reference to specific implementations, it will be recognized that the technology can be implemented in many other, different, forms. To provide a comprehensive disclosure without unduly lengthening the specification, applicants incorporate by reference the patents and patent applications referenced above.

The methods, processes, and systems described above may be implemented in hardware, software or a combination of hardware and software. For example, the auxiliary data encoding processes may be implemented in a programmable computer or a special purpose digital circuit. Similarly, auxiliary data decoding may be implemented in software, firmware, hardware, or combinations of software, firmware and hardware. The methods and processes described above may be implemented in programs executed from a system's memory (a computer readable medium, such as an electronic, optical or magnetic storage device).

The particular combinations of elements and features in the above-detailed embodiments are exemplary only; the interchanging and substitution of these teachings with other teachings in this and the incorporated-by-reference patents/applications are also contemplated.